Amendments to the Specification:

Please replace the paragraph in the specification beginning on page 3, line 15 and ending on page 4, line 7 with the following:

A drawback of the throttling approach described above is that the progress of utility work may be unnecessarily impacted, if the throttling level is too high when the production load [[low]] is low. System administrators have generally competing objectives: (1) ensuring that utilities are throttled sufficiently to prevent undue impact on production work; and (2) ensuring that managed systems are well utilized and/or that utilities finish within an acceptable time. The workload of most commercial systems varies with time, thus it is desirable to change the throttling level in response to changing workloads to optimize system resource allocation. When the production load is high the utilities should be throttled more, but when production load is low the utilities should be throttled less. Moreover, utilities which have little or no impact on the production work may not need to be throttled much or at all.

Please replace the paragraph in the specification beginning on page 13, line 6 and ending on page 13, line 19 with the following:

Referring now to Figure 4 [[5]], the operation of control module 102 is shown in more detail. Preferably, the controller module 102 operates on a periodic basis, i.e. every T seconds, as determined by the system designer. The controller module 102 may be implemented as a proportional-integral (PI) controller. PI controllers are typically robust and offer wide applicability. Any other types of controllers, such as proportional-derivative (PD) or

proportional-integral-derivative (PID) or neural-network based ones, for example, may be used. Also, it is feasible to use controllers that are driven by events other than timer events, for example, by the arrival of specific requests, or based upon having processed a fixed number of requests. The particular choice is specific to the target computer system and is chosen by the system administrator. It is presently preferred that the controller module 102 be implemented as disclosed in the copending and commonly assigned U.S. Patent Application Serial No. 10/427,009, entitled "Adaptive Throttling System for Data Processing Systems", filed on April 30, 2003, and which is hereby incorporated by reference herein.

Please replace the paragraph in the specification beginning on page 15, line 7 and ending on page 15, line 14 with the following:

Figure 5 [[4]] shows a block diagram for an implementation of a throttling system 100 according to another aspect of the present invention in which controller module 102 may obtain and provide information, such as performance metrics, through means other than the manager module 104, e.g., calling operating system interfaces directly. This figure also includes OS interface 110, which may be in operative communication with controller module 102 or manager module 104. Typically, in this embodiment, execution units 106, 108 are not self-throttling and external throttling is accomplished in conjunction with OS interface 110.

Please replace the paragraph in the specification beginning on page 21, line 11 and ending on page 22, line 2 with the following:

The disclosure now turns to a discussion of empirical assessments for the presently preferred SIS throttle. These assessments were made using a modified version of the IBM DB2^{TM} Universal Database v8.1 running on a 4-CPU RS/6000 with 2GB RAM, with the AIX^{TM} 4.3.2 operating system. To emulate client activity, an artificial transaction processing workload was applied which is similar to the industry-standard TPC-C database benchmark. This workload is considered the "production" load. The database is striped over 8 physical disks connected via an serial storage architecture (SSA) disk subsystem. The execution unit (utility) focused on is an online BACKUP of this database. This backup is parallelized, consisting of multiple processes that read from multiple tablespaces, and multiple other processes that write to separate disks.